arXiv:cs/0703117v1 [cs.DC] 23 Mar 2007

Self-adaptive Gossip Policies for Distributed Population-based Algorithms

J.L.J. Laredo[1], E.A. Eiben[2], M. Schoenauer[3], P.A. Castillo[1], A.M. Mora[1], F. Fernández[4], and J.J. Merelo[1]

[1] Department of Architecture and Computer Technology
University of Granada
e-mail: new@geneura.ugr.es
[2] Department of Computer Science
Vrije Universiteit
e-mail: gusz@cs.vu.nl
[3] Projet TAO, INRIA Futurs
LRI, Univ. Paris-Sud
e-mail: marc@lri.fr
[4] Department of Computer Science
University of Extremadura
e-mail: fcofdez@unex.es

Abstract. Gossipping has demonstrate to be an efficient mechanism for spreading information among P2P networks. Within the context of P2P computing, we propose the so-called Evolvable Agent Model for distributed population-based algorithms which uses gossipping as communication policy, and represents every individual as a self-scheduled single thread. The model avoids obsolete nodes in the population by defining a self-adaptive refresh rate which depends on the latency and bandwidth of the network. Such a mechanism balances the migration rate to the congestion of the links pursuing global population coherence. We perform an experimental evaluation of this model on a real parallel system and observe how solution quality and algorithm speed scale with the number of processors with this seamless approach.

1 Introduction

Population-based algorithms are a type of stochastic soft computing techniques widely used as problem-independent solvers in typically NP-hard problems such as graph-search. This paper outlines the general aspects of the development of distributed population-based algorithms in P2P networks with an emphasis on Evolutionary Algorithms (EAs) [7].

EAs are population-based methods with an inherent parallelism that has been widely studied (see e.g. [4] for a survey) and falls mainly under two approaches: master-slave and the island model. In the master-slave mode the algorithm runs on the master and the individuals are sent for evaluation to the slaves, in an approach usually called also *farming*. Using the island model several independent

EAs (islands) are used processing their own population, and exchanging the best individuals between islands with a certain rate [3]. Both cases present major adoption problems in heterogeneous fully decentralized networks such as P2P networks. On one hand, master-slave features do not fit with large-scale system robustness (master represents a single point of failure) and scalability (since it depends on evaluation function cost, and has a bottleneck in the efficiency of the master performing the evolutionary operations). On the other hand, P2P systems do not provide the knowledge of the global environment that the island model needs in order to set parameters such as the number of islands, the population size per island and the migration rate. Some island models also need generations in all nodes to run in lockstep, which calls for homogeneous, synchronized, nodes. This is obviously not the case in P2P ad-hoc networks.

Nevertheless, there is a third, finer grained approach, termed Fully Distributed Model, in which processors host single individuals that evolve on their own. Operations that require more than a single individual (e.g., selection and crossover) take place among a defined set of neighbors (between individuals on different nodes or available locally to a node) [14]. This model is able to adapt to heterogeneous networks since some P2P overlay networks [5] provide a dynamic neighborhood whose size grows logarithmically with respect to the total size of the system in a small-world fashion. Following a gossip style, these small-world networks spread information in an epidemic manner through the whole network (as can be seen in [10, 9]), what means that the risk of having obsolete individuals across the network is minimized as a consequence of the probabilistic global "infection" that the nodes undergo. However, gossiping has to deal with one more question: to maintain the larger coherence in a distributed population among the network, which implies locally to a node not only having high probability of being "infected" but also frequently "infected". We present the approach of the Evolvable Agent model for dealing with such questions.

It is obviously not straightforward to outline a method that takes advantage of those P2P properties, obtaining at the same time high performance and good scalability. That is why we propose a self-adaptive refresh rate over the basis of a gossip scheme which balances the frequency of "infections" to the congestion of links. Within this model, each individual in an evolutionary computation population rises to an autonomous agent by scheduling its own actions.

The main objective of this paper is to provide an empirical assessment of our agent-based evolutionary model which is a step towards a "Fully Distributed Model" for designing EAs in heterogeneous networks. To this end, we perform an experimental evaluation on a real parallel scenario with up to 6 nodes.

The rest of the paper is organized as follows: next (section 2) we describe the state of the art in P2P evolutionary computation and other related subjects. The model is described in section 3, and the particulars of the evolutionary algorithm used here are described in 3.2. Finally, experimental setup is presented in section 4, results in section 5 and some conclusions drawn in 6.

2 Related Work

Due to the diversity of fields that this study involves, it is convenient to revise them in order to set the scope of the work.

Concerning development of P2P distributed computing systems, there are some frameworks such as:

- DREAM [1], which focuses in distributed processing of EAs and uses the P2P network DRM.
- G2DGA [2], equivalent to the previous. It centers on distributed genetic algorithms processing by the use of the network G2P2P.
- JADE (Java Agent Development Framework, available from `http://jade.cselt.it/`), a P2P system which includes agents as software components.

The mentioned DRM is an implementation of the newscast protocol [10]. This protocol has served as a guide for the proposed communication mechanism within this work. Newscast is an epidemic approach where every node shares local information with its neighbourhood by selecting a node from it with uniform probability each certain time (refresh rate). Our communication model is inspired by such a protocol. However, our model considers a dynamic refresh rate which depends on the QoS parameters: latency and bandwidth.

Related to agent-based systems for evolutionary computation, Vacher et al. present in [15] a multiagent approach to solve multiobjective problems. It also describes the implementation of functions and operators of the system. There are some works regarding optimization of parallel evolutionary algorithms; Viveros and Barán [16] propose the combination of parallel evolutionary algorithms with local optimization functions which depends on processor capacities in heterogeneous computational systems. The authors have published related papers on this field: [8] shows that the number of parallel executions must be equivalent to the number of available processors in order to equilibrate computational effort and algorithmic results. [6] report the benefits of considering population size adjustment on runtime. Finally, [11] presents a model (also an agent-based system) where the load of every evolutionary computation experiment is self-adaptive depending on the architecture where it is executed, yielding more efficient results than the classical sequential approach.

In this paper we study an agent-based approach for distributed evolutionary algorithms and propose an asynchronous communication method that allows self-adaptation to different network scenarios and dynamic environments such as P2P systems.

3 Overall Model Description

The overall architecture of our Evolvable Agent Model is depicted on figure 1. It consists of a group of Evolvable Agents (each one running on its own thread) whose main design objective is to carry out the principal steps of evolutionary computation: selection and variation (crossover and mutation). Obviously,

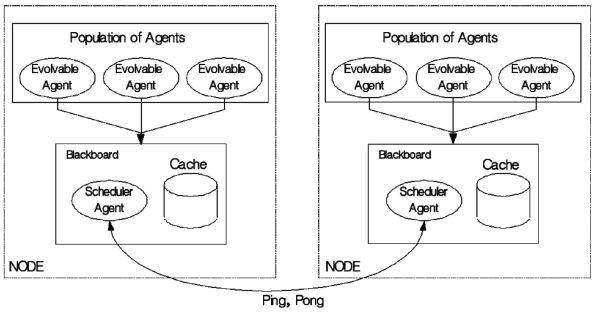

Fig. 1. Overall architecture of the model

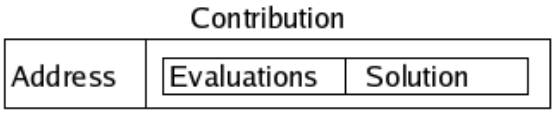

Fig. 2. Format of a cache entry. It provides the following information about a foreign node: Address, number of evaluations performed and one individual of its population termed solution

the key element here is the locally executable selection. Crossover and mutation never involve many individuals, but selection in EAs usually requires a comparison among all individuals in the population. Consider, for example, roulette wheel or rank-based selection.

The agents know the environmental status by means of a blackboard mechanism [13]. The blackboard allows the interchange of information between agents (Agent-Agent) or with cache (Agent-Cache). Furthermore, the blackboard implements a Scheduler Agent that allows information spread among nodes in a gossip style. The messages used among nodes are called *contributions* and their structure matches with a cache entry (figure 2). Thus, instead of the classical view of a population of solutions managed by an oracle, this model proposes a population of agents, each one representing a solution.

3.1 Self-Adaptive Gossip Mechanism

Algorithms 1,2 and 3 show the pseudo-code of the main tasks in the communication process. Each blackboard maintains a cache with a maximum of one entry per node in the network. Each entry follows the contribution format (Figure 2). The cache indexes the entries with the Address field. Therefore, the newest contributions replaces the oldest ones. This process leads the removal of

Algorithm 1 Scheduler Agent

$\Delta T \Leftarrow$ 1 sec.
loop
  sleep ΔT
  Node $\Leftarrow$ Random selected node
  Sol $\Leftarrow$ Selected Solution in P_{agents}
  Contribution $\Leftarrow$ $Num_{Evaluations}$,Sol
  Ping (Node,Contribution)
end loop

Algorithm 2 Ping Handler

Require: Node, Contribution
Cache(Node)$\Leftarrow$ Contribution
Pong(Node,OK)

obsolete individuals and allows a global evolution in a decentralized environment. The scheduling mechanism is carried out by each node as explained next:

- **Algorithm 1** Each ΔT time, a node (the current node) selects another node with uniform probability to establish communication. Current node sends an application level Ping message to the selected node with information about a random solution in the population of agents (P_{agents}) in a contribution format (Figure 2).
- **Algorithm 2** The selected node stores that solution in its cache and sends back an acknowledge message (Pong).
- **Algorithm 3** At the arrival of the Pong, the current node updates its refresh rate (ΔT) with the time spent in the operation.

3.2 Evolvable Agent with Tournament Selection

Algorithm 4 shows the pseudo-code of an Evolvable Agent which uses Tournament Selection. The agent owns a solution (S_t) which it tries to evolve. The selection mechanism works as follows: Each agent selects k (k = tournament size) solutions among other agents' current solutions and solutions stored in cache (which are migrants from network nodes) with uniform probability by means of the blackboard. The two best solutions are stored in "Sols" ready to be recombined by a crossover operator. The crossover returns a single solution S_{t+1} that is mutated and evaluated. If the newly generated solution S_{t+1} is better than the old one S_t, it becomes the current solution. Finally, Blackboard maintains global elitism by storing the best solution found so far in Blackboard.BestSol.

Algorithm 3 Pong Handler

Require: Node
$\Delta T \Leftarrow$ Time used answering the Ping

Algorithm 4 Evolvable Agent with Tournament Selection

```
S_t ⇐ Initialize Agent
Register Agent on the blackboard
loop
  Sols ⇐ Selection(k, Blackboard)
  S_{t+1} ⇐ Recombine(Sols,P_c)
  S_{t+1} ⇐ Mutate(S_{t+1}, P_m)
  S_{t+1} ⇐ Evaluate(S_{t+1})
  if S_{t+1} better than Blackboard.BestSol then
    Blackboard.BestSol ⇐ S_{t+1}
  end if
  if S_{t+1} better than S_t then
    S_t ⇐ S_{t+1}
  end if
end loop
```

4 Experimental Setup

We have carried out an empirical investigation over the Evolvable Agent Model through conducting experiments on a real parallel scenario with up to 6 node.

As a test problem we have chosen the Travelling Salesman Problem (TSP) [12]. The TSP is a classical combinatorial optimization problem widely used to test evolutionary algorithms [17]. In this problem there is a set of $N = 1, \ldots, n$ cities which have to be visited once in such a manner that the path forms a graph cycle that minimizes the travelled distance. We have selected three symmetrical instances with different complexities: *bier127*, *d198* and *lin318*, extracted from TSPLIB[5].

This experiment will provide data on how solution quality (accuracy) and algorithm speed scale with the number of processors. Therefore, we compare results obtained on a single node up to 6 nodes. (Trivial practicalities hindered testing larger networks; further scale-up test are being prepared.) The physical test-bed and the EA main features for the parallel scenario are shown in Table 1. Solution quality is measured by the mean best fitness (MBF) over 30 independent runs. We calculate the speed-up as $S_n = \frac{T_1}{T_n}$, where T_i is the time in seconds spent to reach the termination condition when using i nodes. Linear speed-up is a reference and we use it as the baseline for comparing the scalability results.

Due to the small number of available nodes ($n = 1, \ldots, 6$), we used for all experiments fully connected graph topologies instead of P2P overlay networks. Such scenario grows with a complexity $O(\frac{(n-1)n}{2})$ which intensifies the impact of communication overhead since a real P2P overlay network should grow with a smaller order of complexity in a small-world fashion. As designed, the Scheduler Agent will self-adapt the refresh rate ΔT to the congestion of links.

[5] http://www.iwr.uni-heidelberg.de/groups/comopt/software /TSPLIB95/ Accessed on January 2007

Test-bed	
Number of Nodes	6
Node Processor	2X AMD Athlon(tm)2400+
Network	Gigabit Ethernet
Operating System	Linux 2.6.16-1.2096_FC5smp
Java version	J2RE (build 1.5.0_06-b05)
Evolutionary Algorithm	
Crossover	Order crossover(OX)
Mutation	2-Opt mutation
Termination condition	Max. number of eval.
Parameter values	
Population size	32
Tournament size	7
Number of evaluations	1000000
Probability of crossover	0.7
Probability of mutation	0.1

Table 1. Test-bed and Evolutionary Algorithm settings.

5 Experimental Results

This experimental evaluation focuses on the analysis of the Evolvable Agent model when it scales up to 6 nodes.

Related literature ([8] i.e.) shows how algorithmic results differ in distributed EAs depending either on the availability of computing resources and the spread of the population. A t-Student analysis over the best fitness distributions on the three instances under study reveals that there are no significant differences between most of the them:

- In the bier127 instance, the fitness distributions in 3, 4, 5 and 6 nodes do not present significant differences.
- The same happens in d198 in the case of 1 and 2 nodes, and 4 and 5 nodes.
- In lin318, the fitness distributions in 2, 4, 5 and 6 nodes do not show significant differences.

Since distributed EAs suffer structural changes at population level which modify their algorithmic behaviour, we can conclude from the previous observations that the model under study minimizes the impact of having a distributed population by means of the Scheduler Agent. We have to take into account that the test-bed is composed of a high availability network and the Scheduler Agent adapts migration rates to the network latency and bandwidth. It should be taken into account, too, that differences could be mostly due to the effects of a small population in the algorithm result and might be fixed by using a larger population.

Concerning scalability, figures 3, 4 and 5 represent the speed-up for all problem instances. The linear fit over the data shows a growth close to the baseline (but a bit over it). The data show that the algorithm speed scales well, while

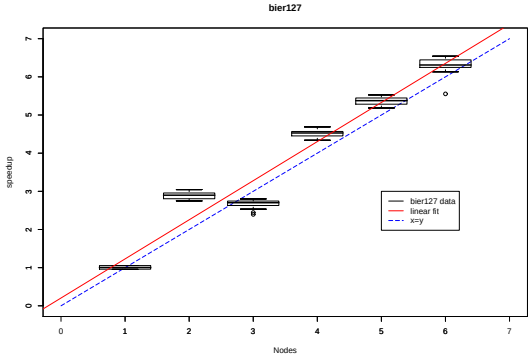

Fig. 3. Speed-up of the Model up to 6 nodes vs. linear speed-up. TSP instance bier127.

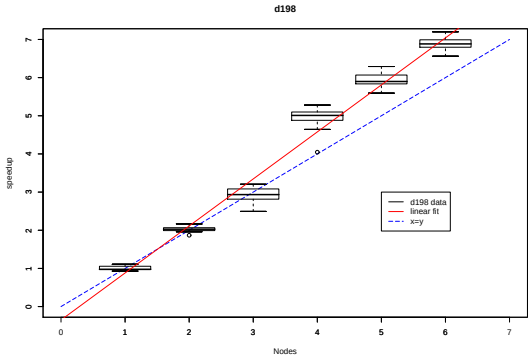

Fig. 4. Speed-up of the model up to 6 nodes vs. linear speed-up. TSP instance d198.

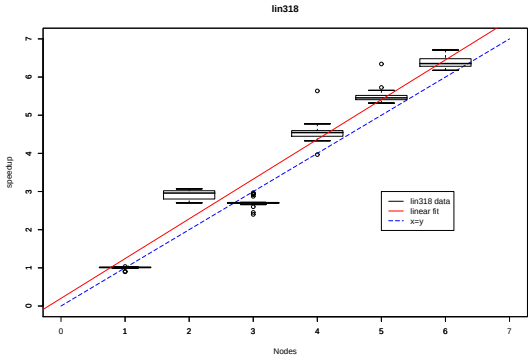

Fig. 5. Speed-up of the model up to 6 nodes vs. linear speed-up. TSP instance lin318.

maintaining solution quality (as has been previously shown). Unfortunately, 6 is still a small network, for extrapolating these results more research is needed. However, within the bounds of this experiment, we consider proved that the Autonomous Agent Model has an efficient scaling behavior, adapting seamlessly to a concurrent as well as a distributed environment.

6 Conclusions and Future Work

In this paper we present an Agent-based approach towards a fully distributed EA model. The model is designed to deal with heterogeneous networks features and specially P2P networks. The evolution process consists in maintaining a population of agents that evolve single solutions. Each agent can access other agents' current solution in operations that needs more than one individual (e.g. selection) by means of the blackboard mechanism described in section 3.

From the proposed experiments we conclude that the model scales with linear gain up to 6 nodes despite the growing complexity in topologies. For a pre-established computational effort, the best fitness distributions in the different test topologies do not reveal significant differences in most of the cases. Therefore, the Evolvable Agent model is a distributed EA model where scalability and quality results are possible both together. Furthermore, such an approach is worth as a proof of concept about self-adaptive gossiping policies for establishing asynchronous migration rates in population-based algorithms.

Future works will have to consider the experimentation in large-scale networks where further conclusions can be reached respecting scalability limitations, adaptation to heterogeneity and algorithmic effects of having high latency links. Within this line we plan to implement the model into a P2P framework such as DREAM [1] which shares its main design objectives.

Acknowledgements

This work has been partially supported under the Project NadeWeb (TIC2003-09481-C04).

References

1. M.G. Arenas, Pierre Collet, A.E. Eiben, M'ark Jelasity, J.J. Merelo, Ben Paechter, Mike Preuss, and Marc Schoenauer. A framework for distributed evolutionary algorithms. In J.J. Merelo Guervós, P. Adamidis, H.G. Beyer, J.L. Fernández-Villaca nas, and H.P. Schwefel, editors, *Parallel Problem Solving from Nature - PPSN VII, 7th International Conference, Granada, Spain, September 7-11, 2002. Proceedings*, number 2439 in Lecture Notes in Computer Science, LNCS, pages 665–675. Springer-Verlag, 2002.
2. Johan Berntsson. G2DGA: an adaptive framework for internet-based distributed genetic algorithms. In *GECCO '05: Proceedings of the 2005 workshops on Genetic and evolutionary computation*, pages 346–349, 2005.

3. Erick Cantú-Paz. Topologies, migration rates, and multi-population parallel genetic algorithms. In Wolfgang Banzhaf, Jason Daida, Agoston E. Eiben, Max H. Garzon, Vasant Honavar, Mark Jakiela, and Robert E. Smith, editors, *Proceedings of the Genetic and Evolutionary Computation Conference*, volume 1, pages 91–98, Orlando, Florida, USA, 13-17 July 1999. Morgan Kaufmann.
4. Erick Cantú-Paz. *Efficient and Accurate Parallel Genetic Algorithms*. Kluwer Academic Publishers, Norwell, MA, USA, 2000.
5. Diego Doval and Donal O'Mahony. Overlay networks: A scalable alternative for P2P. *IEEE Internet Computing*, 7(4):79–82, July/August 2003.
6. A.E. Eiben, E. Marchiori, and V.A. Valkó. Evolutionary algorithms with on-the-fly population size adjustment. In *Parallel Problem Solving from Nature - PPSN VIII*, volume 3242 of *LNCS*, pages 41–50, Birmingham, UK, September 2004. Springer-Verlag.
7. Agoston E. Eiben and J. E. Smith. *Introduction to Evolutionary Computing*. SpringerVerlag, 2003.
8. I. Hidalgo and F. Fernández. Balancing the computation effort in genetic algorithms. In *Evolutionary Computation, 2005. The 2005 IEEE Congress on*, volume 2, pages 1645–1652. IEEE Press, 2005.
9. Márk Jelasity, Alberto Montresor, and Ozalp Babaoglu. Gossip-based aggregation in large dynamic networks. *ACM Trans. Comput. Syst.*, 23(3):219–252, 2005.
10. Márk Jelasity and Maarten van Steen. Large-scale newscast computing on the Internet. Technical Report IR-503, Vrije Universiteit Amsterdam, Department of Computer Science, Amsterdam, The Netherlands, October 2002.
11. J.L.J.Laredo, P.A. Castillo, A.M. Mora, and J.J. Merelo. Estudio preliminar sobre autoadaptación en agentes evolutivos sobre arquitecturas heterogéneas. In *XVII Jornadas de Paralelismo – XVII JP*, pages 389–394, September 2006.
12. E.J. Lawler, J.K. Lenstra, A.H.G. Rinnoy Kan, and D.B. Shmoys. *The travelling salesman problem: A guided tour of combinatorial optimization.* New York: Wiley and Sons, 1985.
13. H. P. Nii. Blackboard systems: The blackboard model of problem solving and the evolution of blackboard architectures. *AI Magazine*, 7(2):38–53, 1986.
14. Marco Tomassini. *Spatially Structured Evolutionary Algorithms: Artificial Evolution in Space and Time*. Natural Computing Series. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2005.
15. Jean-Philippe Vacher, Alain Cardon, Franck Lesage, and Thierry Galinho. Genetic Algorithms in a Multi-Agent System. In *Proceedings IEEE International Joint Symposia on Intelligence and Systems*, pages 17–26, Rockville, MD, USA, 1998.
16. E.R.C. Viveros and B. Barán. Algoritmos genéticos asíncronos combinados para una red heterogénea de computadoras. In *Conferencia Internacional de Ciencia y Tecnología para el Desarrollo*, 1997. http://www.cnc.una.py/cms/invest/download.php?id=46205,66,1.
17. L. Wang, A.A. Maciejewski, H.J. Siegel, and V.P. Roychowdhury. A comparative study of five parallel genetic algorithms using the traveling salesman problem. In *IPPS: 11th International Parallel Processing Symposium*, pages 345–349. IEEE Computer Society Press, 1998.

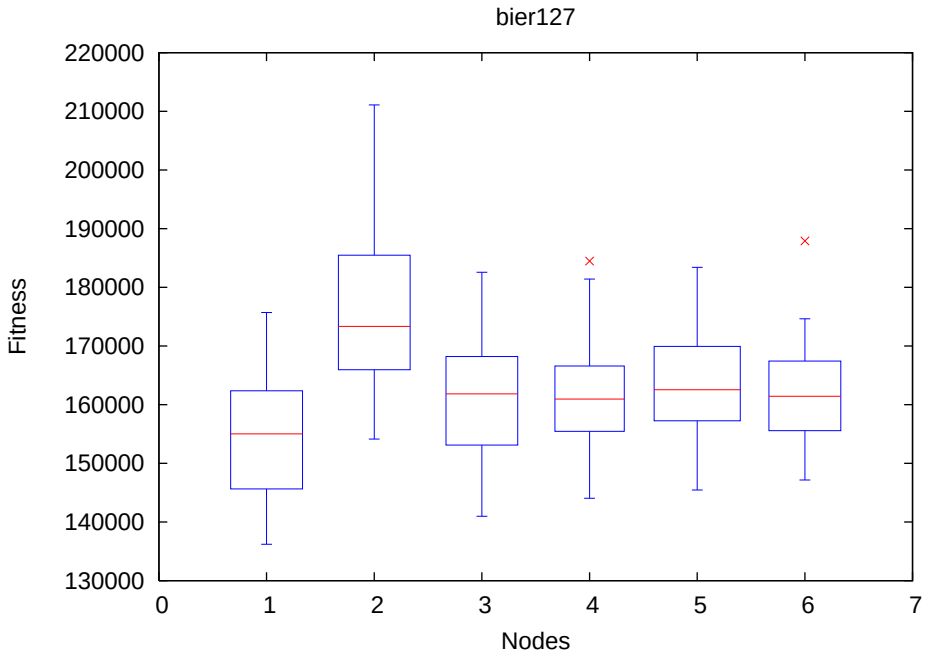
bier127
Fitness
Nodes
220000
210000
200000
190000
180000
170000
160000
150000
140000
130000
0
1
2
3
4
5
6
7

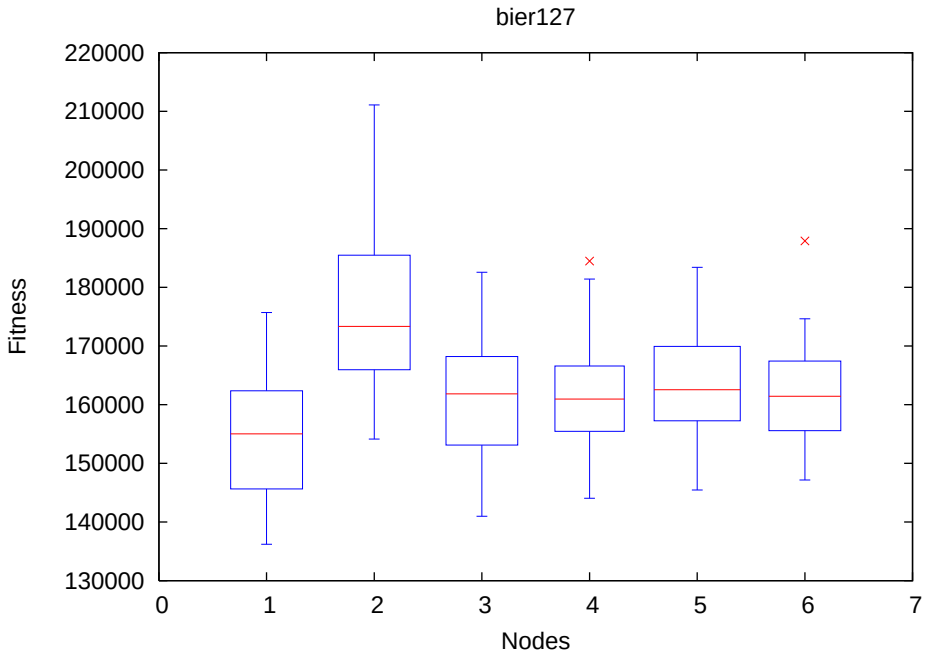
bier127
Fitness
Nodes
220000
210000
200000
190000
180000
170000
160000
150000
140000
130000
0
1
2
3
4
5
6
7

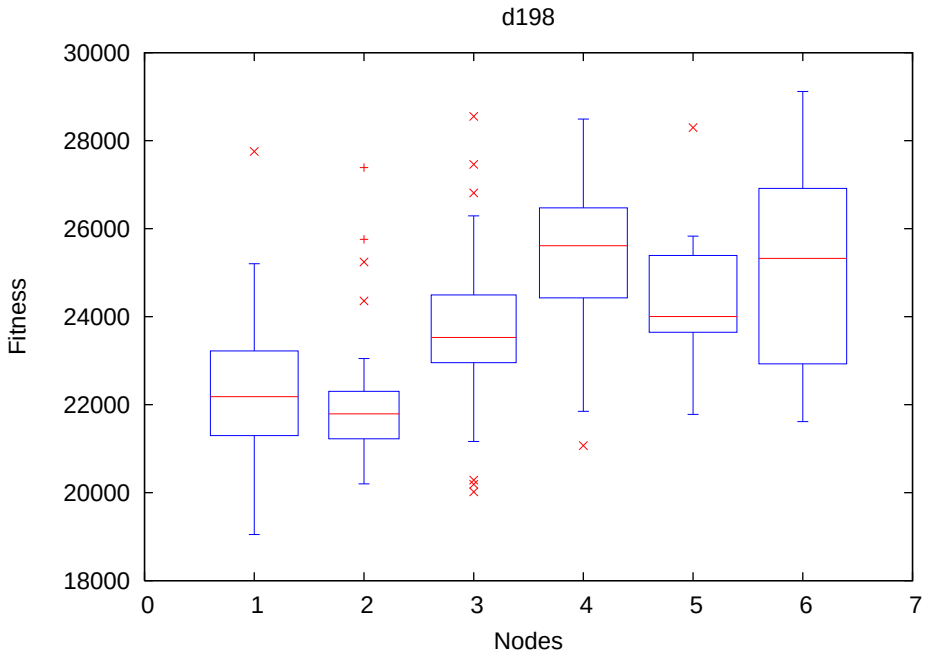
d198
30000
28000
26000
24000
22000
20000
18000
Fitness
0
1
2
3
4
5
6
7
Nodes

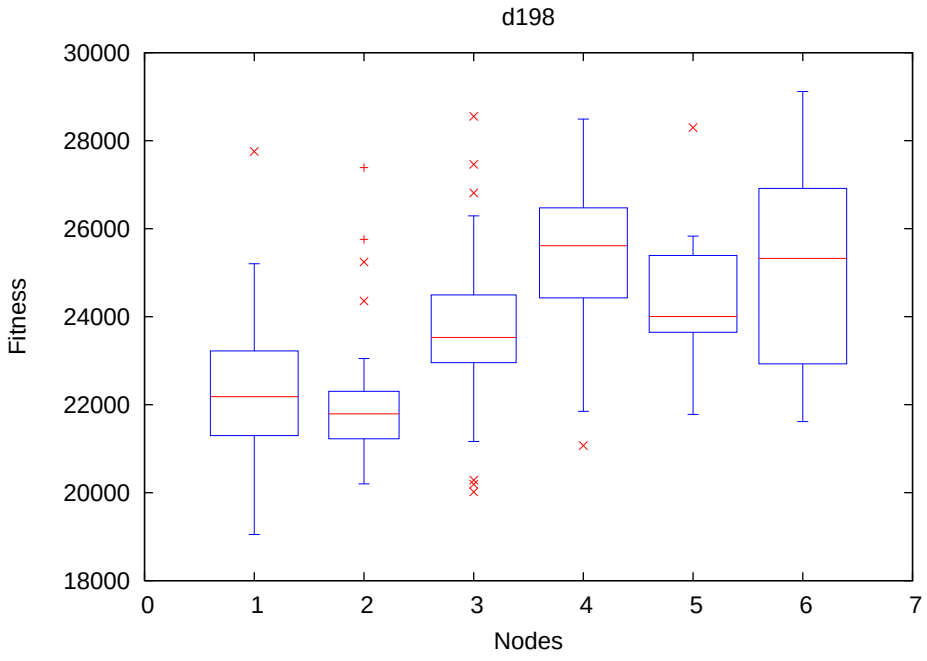
d198
30000
28000
26000
24000
22000
20000
18000
Fitness
0
1
2
3
4
5
6
7
Nodes

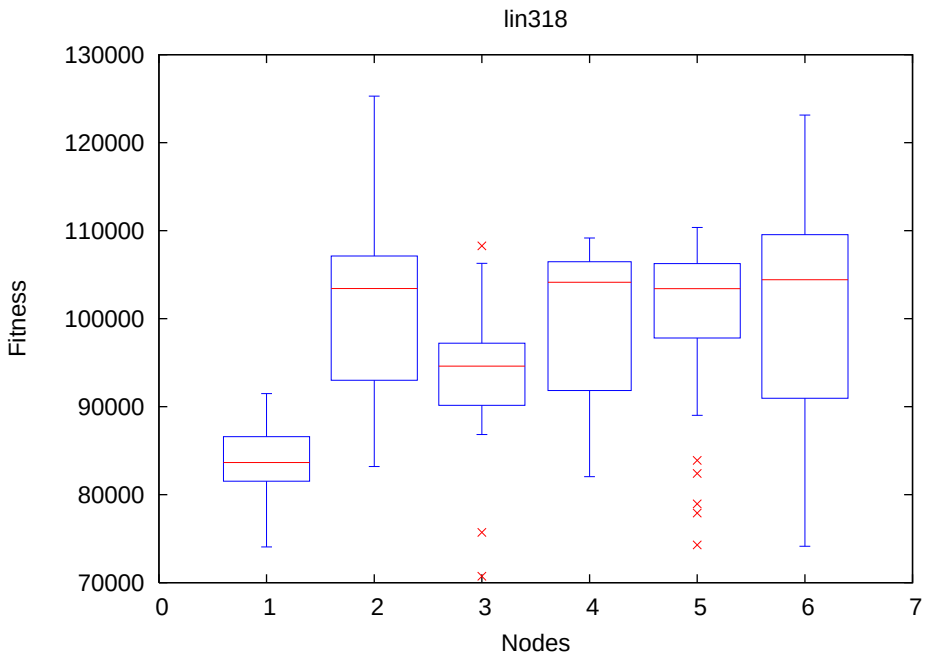
lin318
130000
120000
110000
100000
90000
80000
70000
Fitness
0
1
2
3
4
5
6
7
Nodes

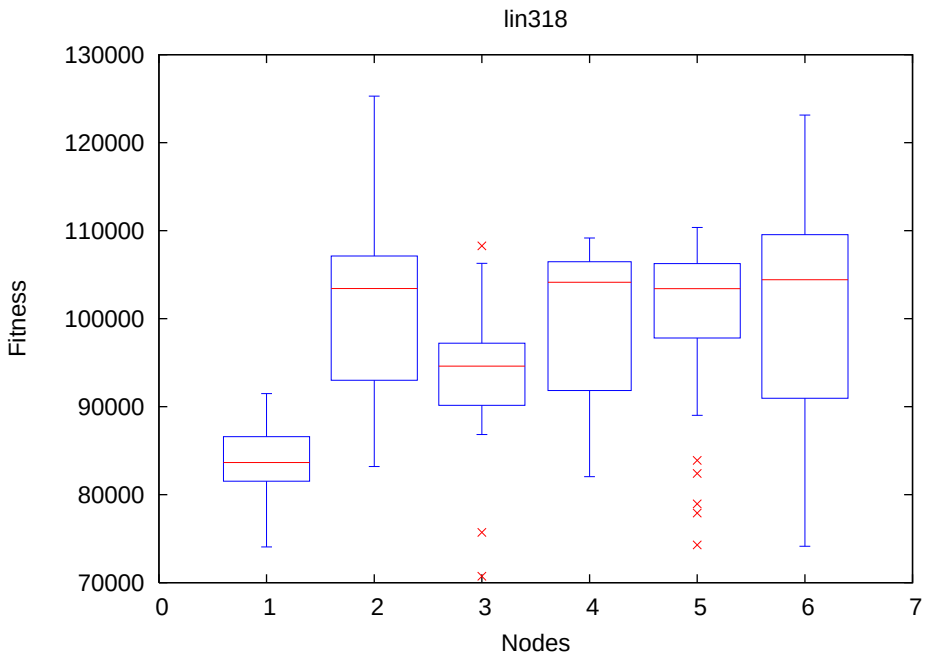
lin318
130000
120000
110000
100000
90000
80000
70000
Fitness
0
1
2
3
4
5
6
7
Nodes

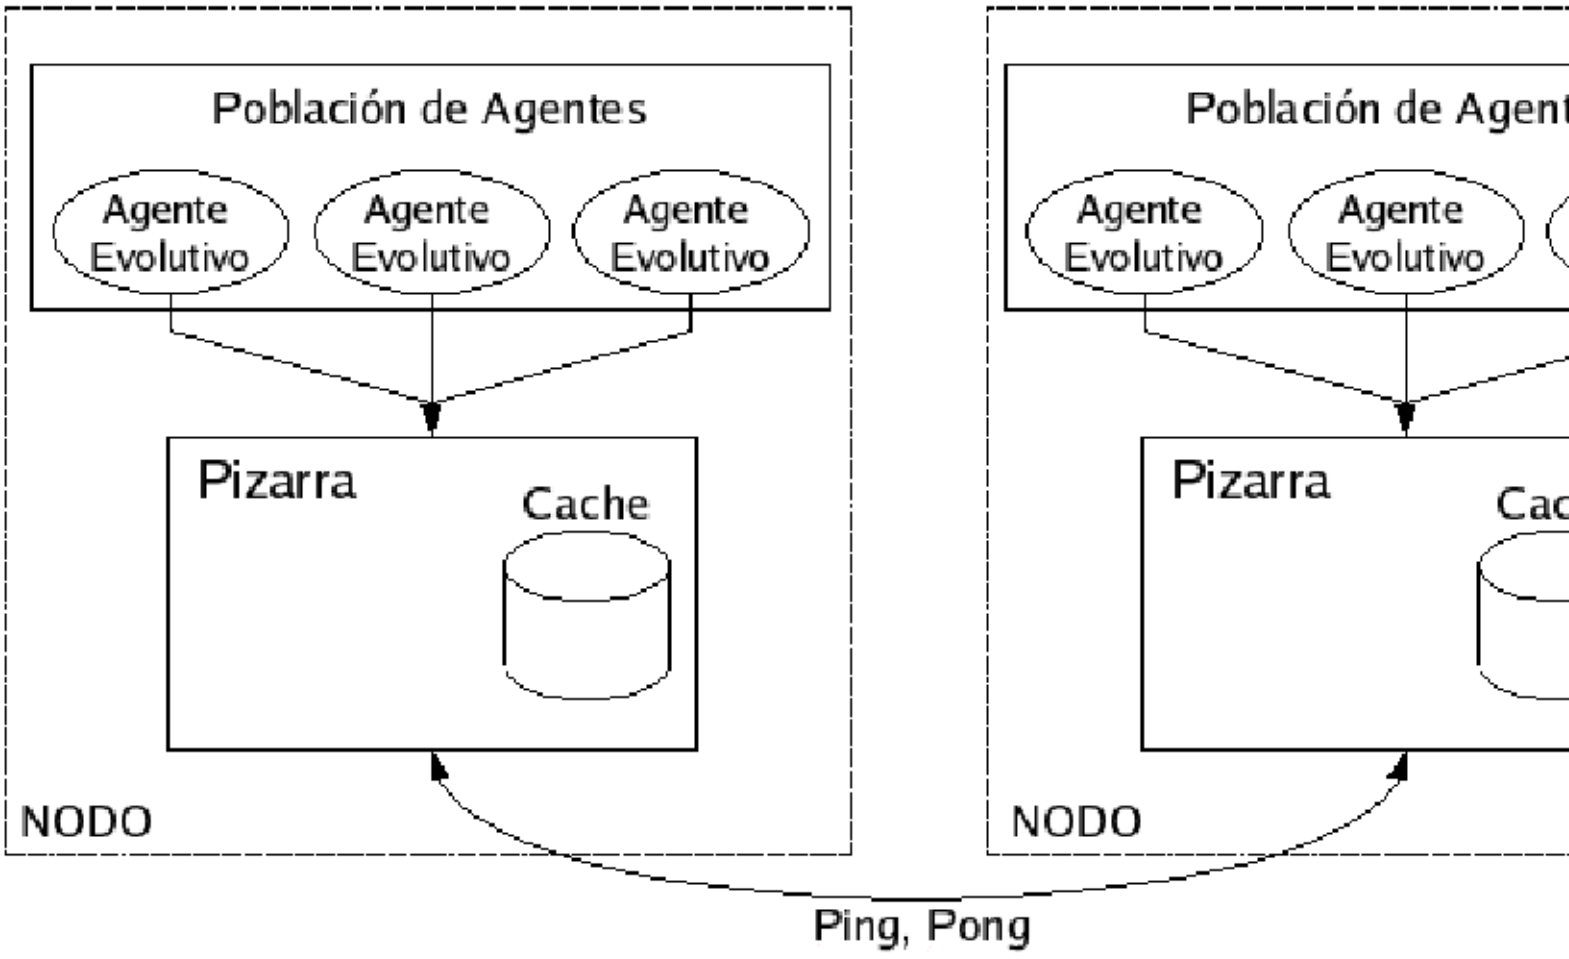
Población de Agentes
Agente Evolutivo
Agente Evolutivo
Agente Evolutivo
Pizarra
Cache
NODO
Población de Agent
Agente Evolutivo
Agente Evolutivo
Pizarra
Cac
NODO
Ping, Pong

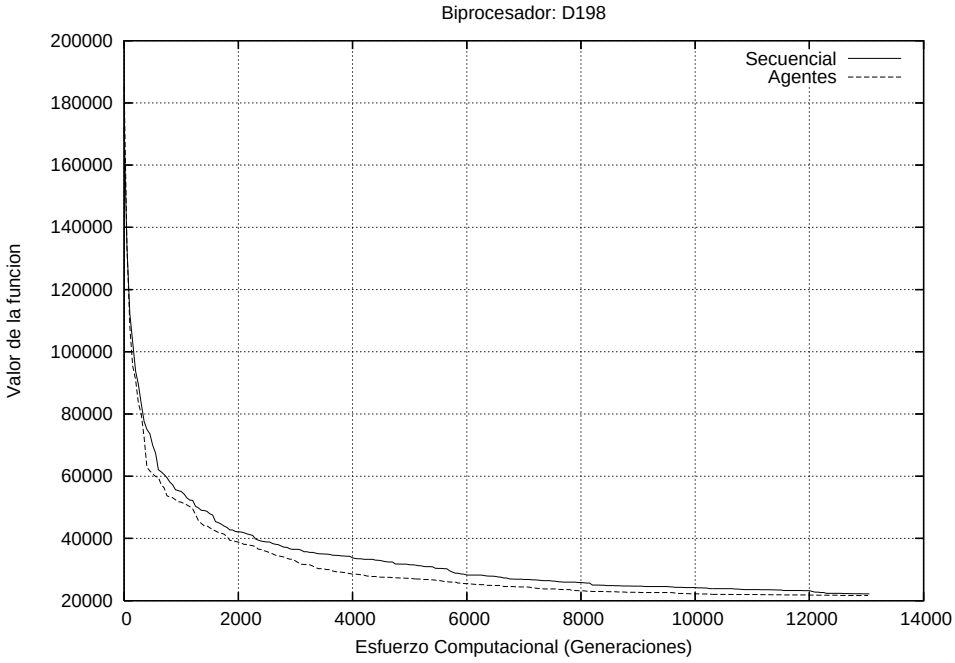
Biprocesador: D198
Secuencial
Agentes
Valor de la funcion
Esfuerzo Computacional (Generaciones)
200000
180000
160000
140000
120000
100000
80000
60000
40000
20000
0
2000
4000
6000
8000
10000
12000
14000

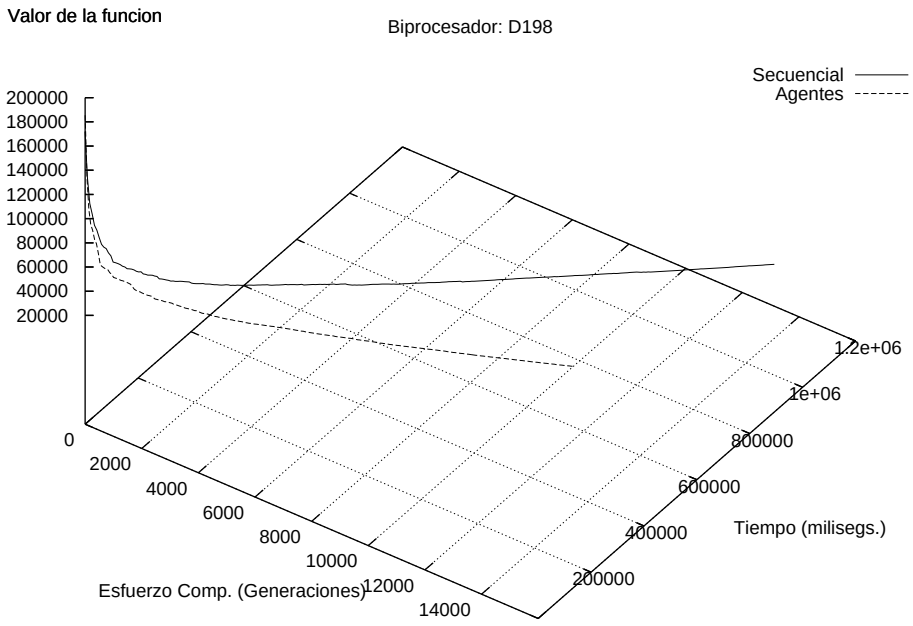
Valor de la funcion
Biprocesador: D198
Secuencial
Agentes
200000
180000
160000
140000
120000
100000
80000
60000
40000
20000
0
2000
4000
6000
8000
10000
12000
14000
Esfuerzo Comp. (Generaciones)
200000
400000
600000
800000
1e+06
1.2e+06
Tiempo (milisegs.)

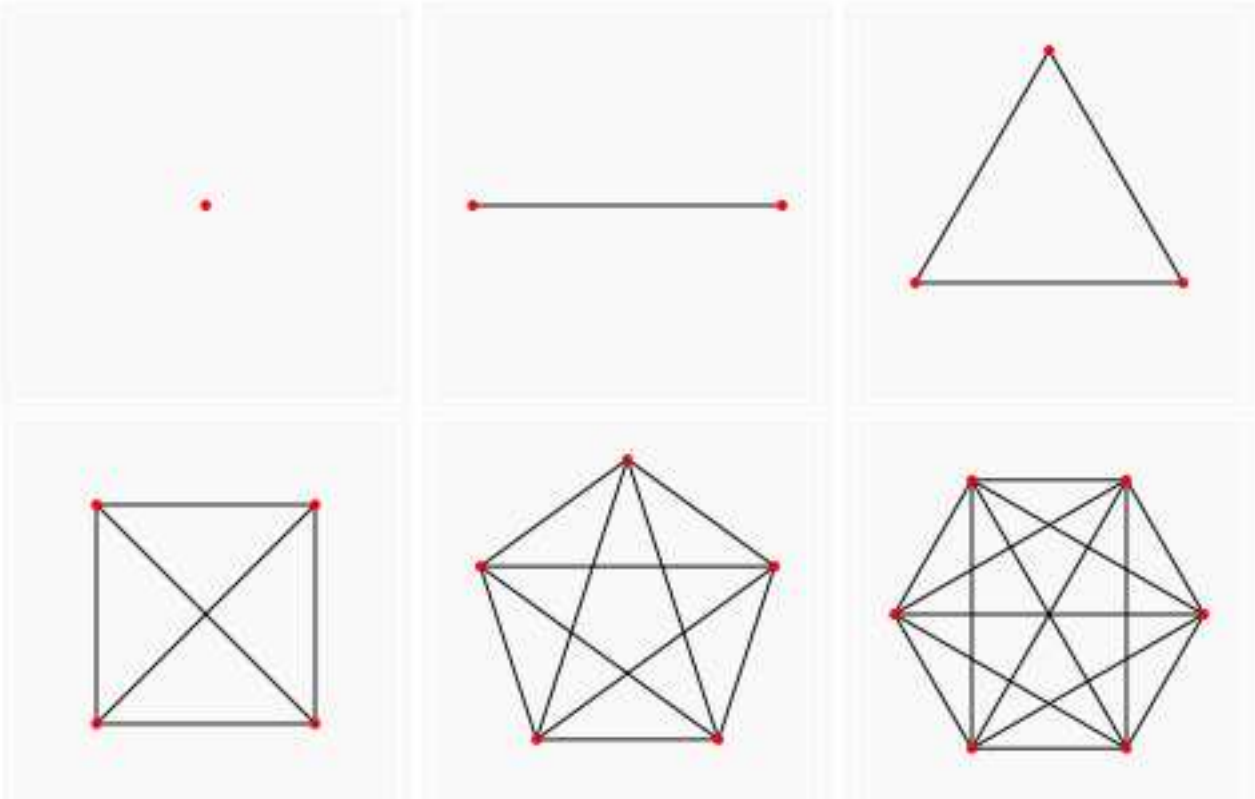

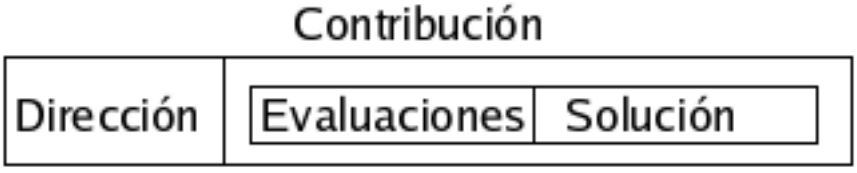
Contribución
Dirección
Evaluaciones
Solución

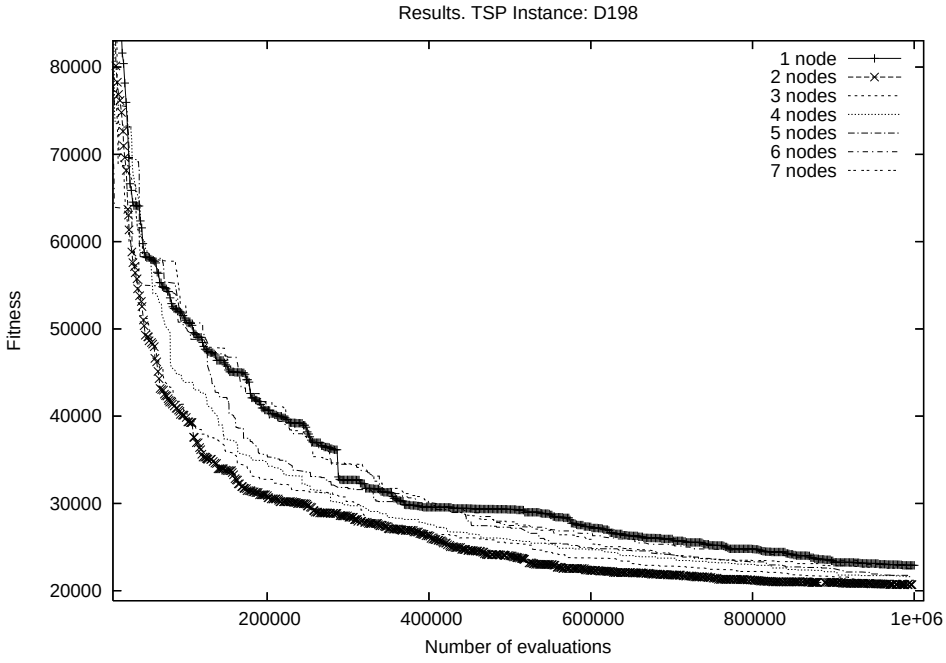
Results. TSP Instance: D198
Fitness
Number of evaluations
80000
70000
60000
50000
40000
30000
20000
200000
400000
600000
800000
1e+06
1 node
2 nodes
3 nodes
4 nodes
5 nodes
6 nodes
7 nodes

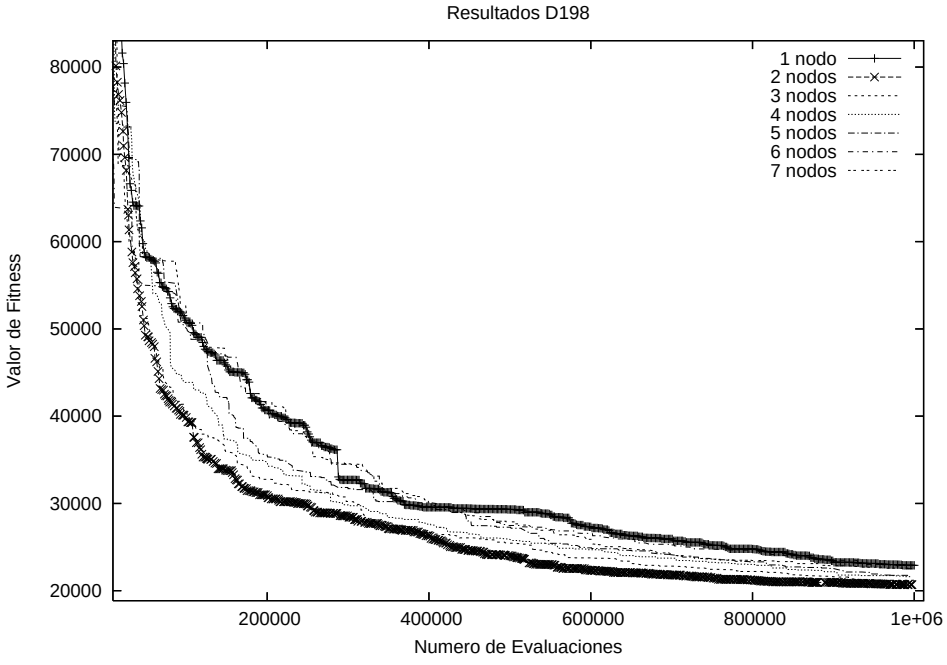
Resultados D198
Valor de Fitness
Numero de Evaluaciones
80000
70000
60000
50000
40000
30000
20000
200000
400000
600000
800000
1e+06
1 nodo
2 nodos
3 nodos
4 nodos
5 nodos
6 nodos
7 nodos

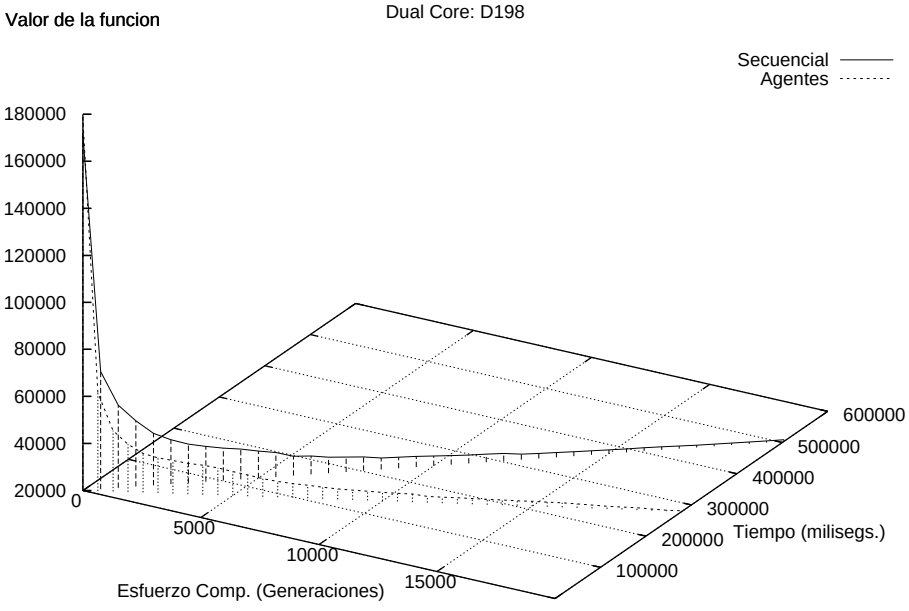
Valor de la funcion
Dual Core: D198
Secuencial
Agentes
180000
160000
140000
120000
100000
80000
60000
40000
20000
0
5000
10000
15000
Esfuerzo Comp. (Generaciones)
600000
500000
400000
300000
200000
100000
Tiempo (milisegs.)

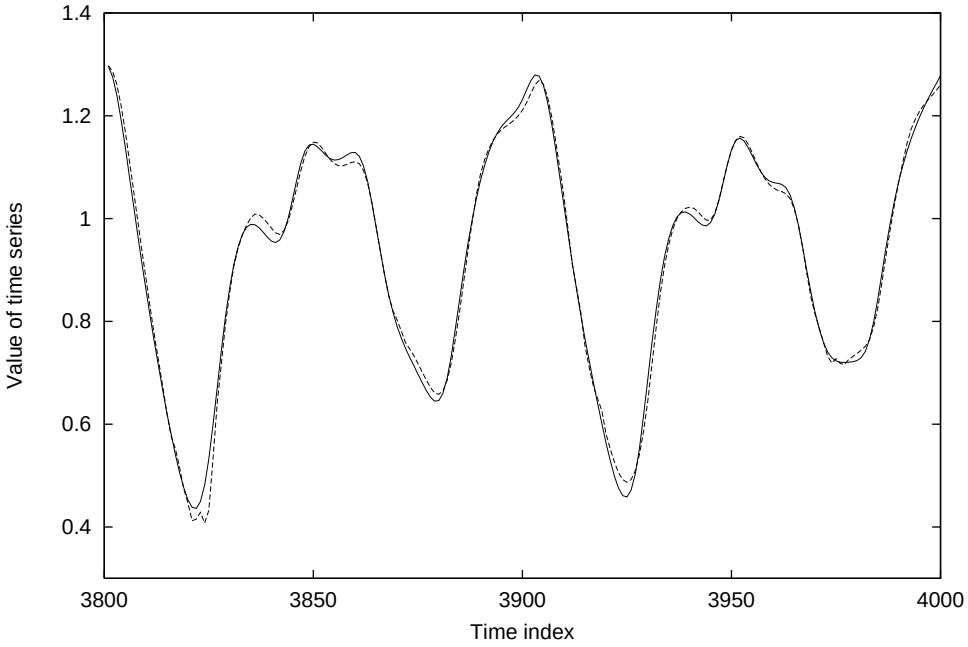
1.4
1.2
1
0.8
0.6
0.4
Value of time series
3800
3850
3900
3950
4000
Time index

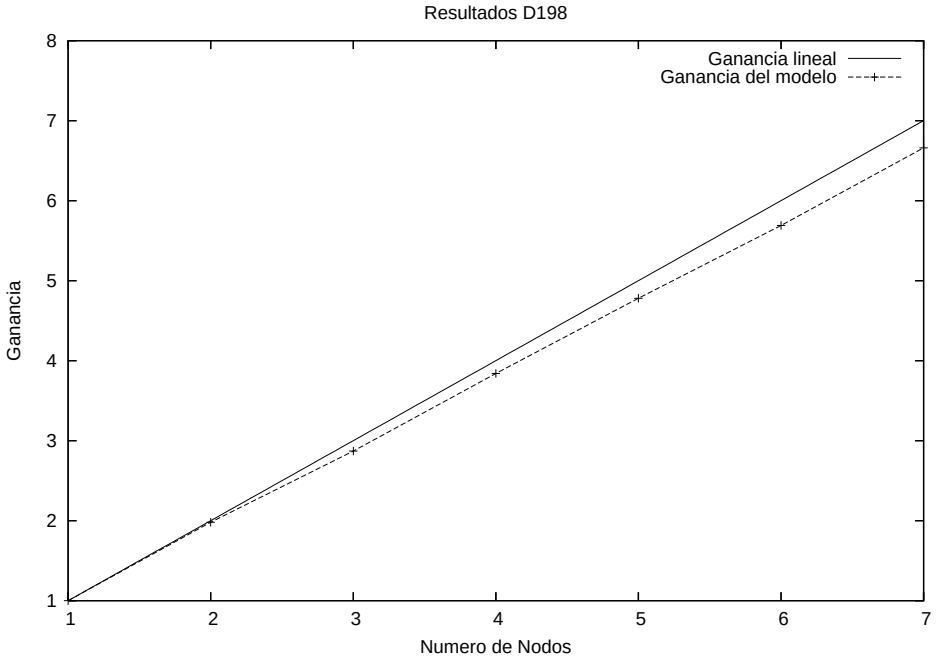
Resultados D198
Ganancia lineal
Ganancia del modelo
Ganancia
Numero de Nodos
1
2
3
4
5
6
7
8

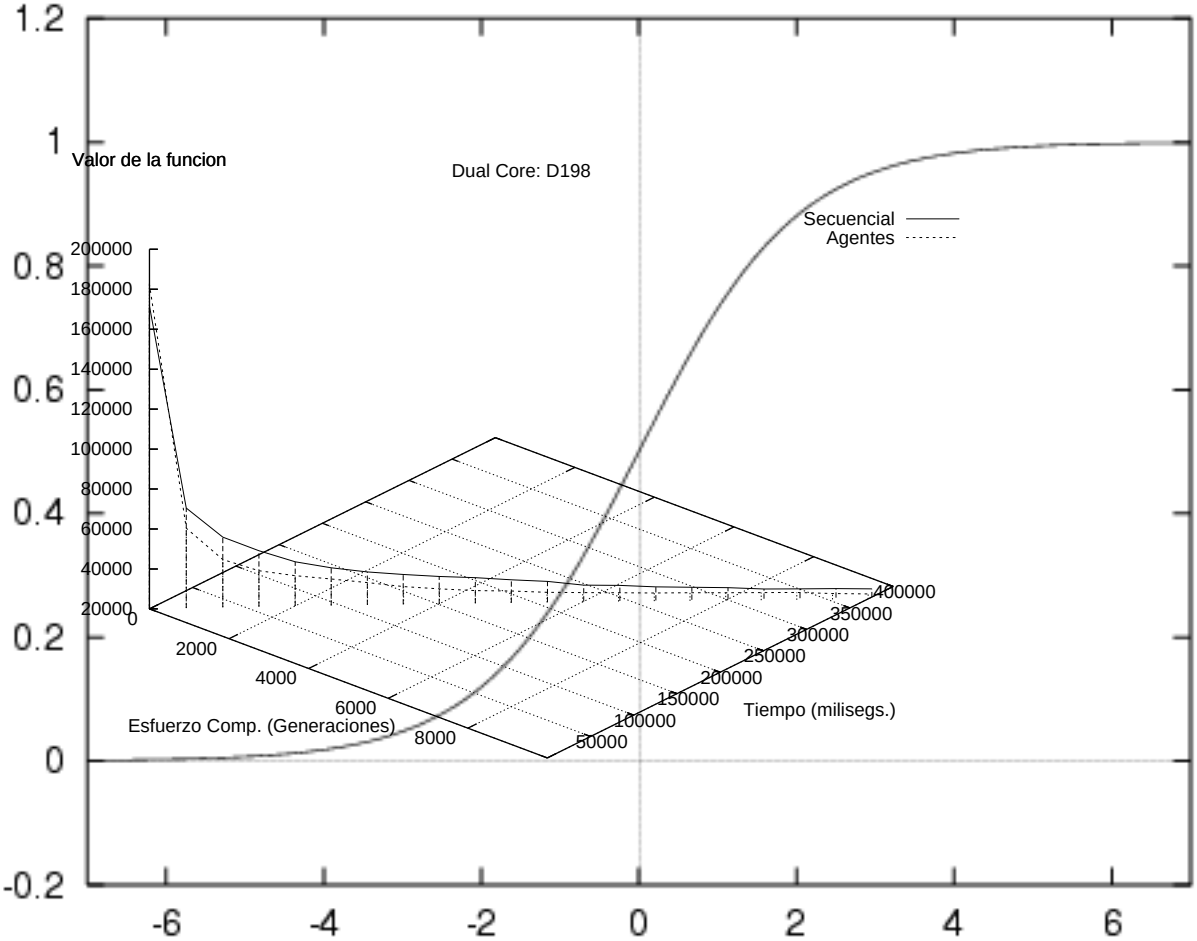
Valor de la funcion
Dual Core: D198
Secuencial
Agentes
Esfuerzo Comp. (Generaciones)
Tiempo (milisegs.)
1.2
1
0.8
0.6
0.4
0.2
0
-0.2
-6
-4
-2
0
2
4
6
200000
180000
160000
140000
120000
100000
80000
60000
40000
20000
0
2000
4000
6000
8000
400000
350000
300000
250000
200000
150000
100000
50000

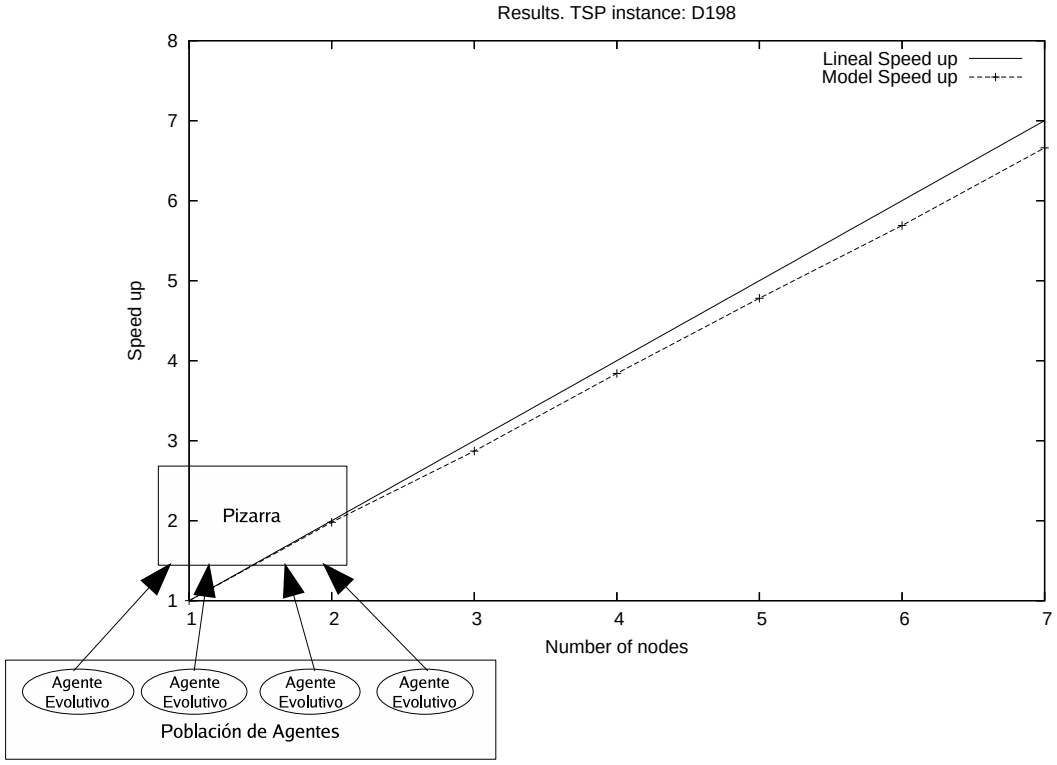
Results. TSP instance: D198
Lineal Speed up
Model Speed up
Speed up
Number of nodes
Pizarra
Agente Evolutivo
Agente Evolutivo
Agente Evolutivo
Agente Evolutivo
Población de Agentes